

Triadic Consensus

A Randomized Algorithm for Voting in a Crowd

Ashish Goel¹ and David Lee¹

Stanford University {ashishg, davidtlee}@stanford.edu

Abstract. Typical voting rules do not work well in settings with many candidates. If there are just several hundred candidates, then even a simple task such as choosing a top candidate becomes impractical. Motivated by the hope of developing group consensus mechanisms over the internet, where the numbers of candidates could easily number in the thousands, we study an urn-based voting rule where each participant acts as a voter and a candidate. We prove that when participants lie in a one-dimensional space, this voting protocol finds a $(1 - \epsilon/\sqrt{n})$ approximation of the Condorcet winner with high probability while only requiring an expected $O(\frac{1}{\epsilon^2} \log^2 \frac{n}{\epsilon^2})$ comparisons on average per voter. Moreover, this voting protocol is shown to have a quasi-truthful Nash equilibrium: namely, a Nash equilibrium exists which may not be truthful, but produces a winner with the same probability distribution as that of the truthful strategy.

1 Introduction

Voting is often used as a method for achieving consensus among a group of individuals. This may happen, for example, when a committee chooses a representative or friends go out to watch a movie. When the group is small, this process is relatively easy; however, for larger groups, the typical requirement of ranking all candidates becomes impractical and heuristics are often applied to narrow down opinions to a few representative ones before a vote is taken.

This problem of large-scale preference aggregation is even more interesting in light of the rising potential of crowdsourcing. Suppose that a city government wanted to ask its constituencies to contribute solutions for an “ideal budget that cuts 50 percent of the deficit”.¹ Soliciting such proposals may be relatively straightforward; however, it is not clear how these proposals should be aggregated. In particular, a participant cannot even look through each proposal, making seemingly simple tasks such as choosing top ranked proposals, difficult. A solution to this problem would enable a new level of collaboration, a key step towards unleashing the full potential of crowdsourcing.

In this paper, we propose a randomized voting rule designed for scenarios like the above. In our problem setting, each participant submits exactly one proposal, representing his or her stance on the question of interest. A random

¹ See, for example, widescope.stanford.edu, aimed at tackling the federal budget deficit

triad of participants is then selected and each selected member is made to vote between the other two. Roughly speaking (details are elaborated in Sect. 2.1), if there is a three-way tie, the participants are thrown out from the election; otherwise, the losers are replaced by ‘copies’ of the winner. This is then repeated until there is a single participant remaining, who is declared the winner.

We show that for single peaked preferences, Triadic Consensus converges approximately to the Condorcet winner² with high probability, while only requiring an average of $\sim \log^2 n$ (conjectured to be $\sim \log n$) comparisons per individual. As an illustration, recall our motivating scenario of a city government crowdsourcing a question to its constituencies. Suppose that the city finds ten thousand participants and that the participant preferences are single-peaked. Clearly, looking through all 10000 proposals to perform the simple task of voting for a favorite is impossible. With Triadic Consensus, each participant would make an average of 177 (conjectured to be 13.3) pairwise comparisons for the algorithm to produce a winner. This winner would be between proposals 4950 and 5050 with 95 percent probability and between proposals 4900 and 5100 with 99.99 percent probability. In other words, the winner would be in the top 1 percent of submitted proposals with 95 percent probability and in the top 2 percent of submitted proposals with 99.99 percent probability.

In addition, we show that Triadic Consensus has nice properties for protecting against manipulation. Suppose that the rankings of candidates are induced from an underlying distance metric and suppose that each candidate has a concave utility in that distance. Then Triadic Consensus has a *quasi-truthful* Nash equilibrium. Specifically, (see Sect. 2.2) a Nash equilibrium exists which may not be truthful, but still chooses a winner *with the same probability distribution* as if every participant voted truthfully. Surprisingly, we achieve this result by counterintuitively allowing voters to express cyclical preferences (e.g. $a > b$, $b > c$, and $c > a$).

1.1 Related Work and Our Contributions

Given the long history of work on voting theory, it is not surprising that the problems we tackle have been, for the most part, thought about before. Here, we give a brief overview of related work, followed by a summary of our contributions. For in-depth reading, we refer the reader to Brandt et al. [1].

Voting rule criteria One of the earliest criteria introduced for evaluating voting rules is known as the Condorcet criteria, introduced by Marquis de Condorcet³. It states that if a candidate exists who would win against every other candidate in a majority election, then this candidate should be elected. Unfortunately, such a candidate does not always exist. Since then, many other criteria

² The candidate who would beat any other candidate in a pairwise majority election.

In single dimensional spaces, this happens to be the median participant.

³ See Young [2] for a fascinating historical description of the early work of Condorcet.

have been introduced as ways to evaluate voting rules. However, in the surprising result known as Arrow's Impossibility Theorem, Arrow [3] proved that there were three desirable criterion that no deterministic voting rule could satisfy. This was expanded by Pattanaik and Peleg [4] to show that a similar result holds for probabilistic voting rules.

Strategic manipulation This sparked a wave of impossibility results, including the classical Gibbard-Satterthwaite Impossibility Theorem. Define a voting rule to be *strategy-proof* if it is always in a voter's interest to submit his true preference, regardless of the other voter rankings. Gibbard [5] and Satterthwaite [6] independently showed that all deterministic, strategy-proof voting rules must either be dictatorships or never allow certain candidates to win. This was extended to show that only very simple probabilistic voting rules were strategy-proof[7].

Numerous attempts at circumventing these impossibility result have been made. Bartholdi et al. [8] first proposed using computational hardness as a barrier against manipulation in elections. However, despite many NP-hardness results on manipulation of voting rules[9], it was shown that there do not exist any voting rules that are *usually* hard to manipulate[10].

Procaccia [11] used the simple probabilistic voting rules of Gibbard [7] to approximate common voting rules in a strategy-proof way, but the approximations are weak and they show that, for many of these voting rules, no strategy-proof approximations can be much stronger. Birrell and Pass [12] extended this idea to approximately strategy-proof voting, proving that there exist tight approximations of any voting rule that are close to strategy-proof. Recently, Alon et al. [?] studied the special case of approval voting when voters and candidates coincide. They show that even though no deterministic strategy-proof mechanism has a finite approximation ratio, a randomized strategy-proof mechanism exists which has a good approximation ratio.

Communication complexity When the number of candidates is large, it is important to study voting rules from the perspective of the burden on voters. Conitzer and Sandholm [13] studied the worst case number of bits that voters need to communicate (e.g. pairwise comparisons) in order to determine the ranking or winner of common voting rules; for many of these voting rules, it was shown that the number of bits required is essentially the same as what is required for reporting the entire ranking. In addition, they showed [14] that for many common voting rules, determining how to elicit preferences efficiently is NP-complete, even when perfect knowledge about voter preferences is assumed. Lu and Boutilier [15] proposed the idea of reducing communication complexity under approximate winner determination. Though they do not present theoretical guarantees, they propose a regret minimizing algorithm and show significant reductions in communication when run on experimental data sets.

Single-peaked preferences One special case that avoids the many discouraging results above is that of single-peaked preferences[16] (or other domain

restrictions). Single-peaked preferences are those for which candidates can be described as lying on a line. Every voter’s utility function is peaked at one candidate and drops off on either side. For such preferences, a Condorcet winner always exists and is the candidate who is the median of all voter peaks. This winner can be found by the classical median voting rule, which has each voter state their peak and returns the median of these peaks. It turns out that the median voting rule is both strategy-proof[17] and has a low communication complexity of $O(n \log m)$ [18], where n is the number of voters and m is the number of candidates. Conitzer [19] also studies the problem of eliciting voter preferences or the aggregate ranking using comparison queries.

The median voting rule has one weakness: it requires knowledge of an axis, which can make it impractical in practice. First, the *algorithm* requires knowledge of the axis in order to pick the median of peaks. When an axis isn’t known, Escoffier et al. [18] provides an $O(mn)$ algorithm for finding such an axis with additional queries, but with no strategic guarantees. Second, the *voter* also requires knowledge of the axis. In situations where proposals have multiple criterion, but are still single peaked (for example, in a linear combination of the criterion), it may not be obvious to the voter where the axis is. Third, and more subtle, even if an axis is known, it may not be practical to *express* a voter’s position on this axis. Take, for example, the canonical liberal-conservative axis used to support the single-peaked setting. It is obvious that one extreme of the axis is an absolute liberal and that the other is an absolute conservative. But how would a voter express any position in between? It would not make sense for a voter to express his or her peak as “seventy percent liberal”.⁴

Our contributions Triadic Consensus solves the previous problems by eliminating the need for an axis. The only task voters are required to perform is a series of comparisons between two candidates. Likewise, the central algorithm does not require any knowledge about proposal positions. With these properties, we prove the following guarantees (as made precise in Sects. 3 and 4):

1. For single-peaked preferences, Triadic Consensus finds a $(1 - \epsilon/\sqrt{n})$ approximation of the Condorcet winner with high probability with a communication complexity of $O(\frac{n}{\epsilon^2} \log^2 \frac{n}{\epsilon^2})$, i.e. $\sim n \log^2 n$ (conjectured to be $\sim n \log n$) for a $1 - \frac{1}{\sqrt{n}}$ approximation and $\sim \frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon^2}$ for a $1 - \epsilon$ approximation.
2. For a single-dimensional setting, Triadic Consensus has a quasi-truthful Nash equilibrium when participants have concave utility functions.

These results are especially interesting given that they are coupled with the following novel concepts:

1. *A localized consensus mechanism for large groups.* We propose Triadic Consensus as an approach for large groups to make decisions using small decentralized decisions among groups of three.

⁴ Note that he cannot just state his favorite candidate as his peak because this would require looking through all n candidates.

2. *Quasi-truthful voting rules and cyclical preferences.* When each participant is a voter and a candidate, we demonstrate that allowing participants to express cyclical preferences ($a > b$, $b > c$, and $c > a$) can introduce strategies that detect and protect against strategic manipulation.

Outline of the paper Before continuing, we describe the structure of the remaining sections. In Sect. 2, we detail Triadic Consensus and introduce the notion of quasi-truthfulness. This is followed by Sect. 3, which presents the approximation and communication complexity results, and Sect. 4, which describes the quasi-truthfulness results. Finally, Sect. 5 concludes with future directions.

2 Triadic Consensus and Quasi-truthfulness

2.1 Triadic Consensus

ALGORITHM 1: Triadic Consensus

Input: An urn with k labeled balls for each participant $1, 2, \dots, n$

Output: A winning candidate i .

while *there is more than one label* **do**

Sample three balls (with labels x, y, z) uniformly at random with replacement;

$w = \text{TriadicVote}(x, y, z)$;

if $w \neq \emptyset$ **then**

Relabel all the sampled balls with the winning label w ;

else

/* For example, remove the three sampled balls from the urn */
TriadicMechanism(x, y, z);

if *at least one ball remains* **then**

return the id of any remaining ball;

else

return the id of a random ball from the last removed;

ALGORITHM 2: TriadicVote

Input: Candidates x, y, z

Output: One of $\{x, y, z\}$ if there is a winner, \emptyset otherwise

if *two of more of x, y, z have the same id* **then**

return the majority candidate;

x votes between y and z ; y votes between x and z ; z votes between x and y ;

if *each received exactly one vote* **then**

return \emptyset ;

else

return the candidate with two votes;

Triadic Consensus applies to scenarios where the set of candidates and voters coincide. We use x to refer to both the participant x and the candidate solution

ALGORITHM 3: The Remove mechanism

Input: Balls x , y , and z

Remove the three sampled balls from the urn;

ALGORITHM 4: The RepeatThenRemove mechanism

Input: Balls x , y , and z

$w = \text{TriadicVote}(x, y, z)$;

if $w \neq \emptyset$ **then**

 | Relabel all the sampled balls with the winning label w ;

else

 | Remove(x, y, z);

that he or she proposes. For $x, y, z \in \{1, 2, \dots, n\}$, we use \succ_x to denote the ranking of participant x and $y \succ_x z$ to denote that x prefers y over z .⁵

The best way to understand Triadic Consensus (Alg. 1) is to imagine an urn with balls, each of which is labeled by a participant id. The urn starts with k balls for each of the n participants.⁶ At each step, the algorithm samples three balls uniformly at random (with replacement) and performs a TriadicVote (Alg. 2) on the three corresponding participants.

If the three participants x , y , and z are unique, the TriadicVote subroutine consists of a single comparison for each of the selected participants: x votes between y and z , y between x and z , and z between x and y . These votes can be distributed in some permutation of 2, 1, 0 or split 1, 1, 1. In the first case, the participant who received two votes is returned as the winner. In the second case, a tie (represented as \emptyset) is returned. If two or more of the selected ids are the same, i.e. are the same person, then he is automatically returned as the winner.

If a winner was returned from the TriadicVote, then the three balls are relabeled with the winning id and placed back into the urn; otherwise, one of several mechanisms can be applied to resolve the tie. This process is repeated until there is only one participant id remaining, which is declared the winner. A helpful intuition is that the mechanism implemented in the case of a tie can be thought of as a deterrent for manipulation since a three-way tie can never occur if participants are voting truthfully (see Sect. 2.2).

In our paper, we propose two possible mechanisms, each of which has a quasi-truthful Nash equilibrium. The simplest is Remove (Alg. 3), in which the three balls are simply removed. In RepeatThenRemove (Alg. 4), the three balls are made to vote again; if there is another three way tie, then they are removed. Surprisingly, repeating the TriadicVote before elimination results in a simpler (and more practical) strategy that is a quasi-truthful Nash equilibrium.

⁵ We assume a strict ordering, but it is not hard to generalize the algorithm to ties.

⁶ The intuition for k is that it is a tradeoff between approximation and time. Increasing k makes the approximation tighter, but requires more comparisons to converge.

2.2 Truthfulness and Quasi-truthfulness

For our analysis of strategic behavior, we will assume that each individual is represented as a point x in some space X and that his or her preference ranking is induced by a distance metric $d(x, \cdot)$ on X . If $d(x, y) \leq d(x, z)$, then $y \succ_x z$; that is, x prefers proposals that are closer to him. Since the individuals voting in a TriadicVote are also the candidates being voted for, there can never be a three-way tie in a truthful vote. Otherwise, all three of $d(x, y) < d(x, z)$, $d(y, z) < d(y, x)$, and $d(z, x) < d(z, y)$ must be simultaneously true, which is impossible so long as $d(\cdot, \cdot)$ satisfies the natural property that $d(x, y) = d(y, x)$.⁷

Consider a TriadicVote between participants x , y , and z . If they vote truthfully, then there are situations when players may be incentivized to deviate as in Ex. 1.

Example 1. Four participants lie in space $X = \mathbb{R}$ at positions 0, 5, 6, and 7. Suppose participants 0, 5, and 7 are selected for a TriadicVote. Since they are voting truthfully, 0 votes for 5, 5 votes for 7, and 7 votes for 5. As a result, 5 wins and the resulting urn consists of three balls for 5 and one for 6.

Now suppose participant 7 were to vote strategically for participant 0. This would result in a tie and all the selected balls would be eliminated, leaving only participant 6. Clearly, participant 7 would prefer this second situation.

At this point, we might note that the truthful winner's vote (e.g. 5) does not change the result and that he can use his vote to disincentivize others from manipulating the TriadicVote. We define any such behavior to be quasi-truthful when it results in the same outcome as that of truthful voting.

Example 2. Suppose that the participants of Ex. 1 are trying to minimize the expected distance of the winning proposal to their position. Then a quasi-truthful strategy would be for 0 to vote for 5, 5 to vote for 0 and 7 to vote for 5. As in Ex. 1, 5 wins and the resulting urn consists of three balls for 5 and one for 6.

Now suppose participant 7 deviates from this strategy and votes for 0. Then participant 0 gets two votes and he wins. The resulting urn consists of three balls for 0 and one for 6, which is clearly worse for participant 7. Likewise, suppose participant 0 deviates from this strategy and votes for 7. Then there is a three-way tie and all selected balls get eliminated. The resulting urn consists of a single ball for 6, which is clearly worse for participant 0.

From this example, we get the intuition for Alg. 5, a quasi-truthful Nash for the Remove mechanism. If a participant (y WLOG) is the truthful winner, then he should look for the participant who would prefer a win for him over a removal of all three balls (x WLOG). Then if y votes for x , this will disincentivize x from deviating. Since z can only cause x to win by deviating, it seems intuitively bad for him to deviate since the winner will be strictly farther. If all players

⁷ If voter rankings allow ties, then a three-way tie could be truthful. In this case, a simple generalization of the TriadicVote that aligns with the intuition of punishing manipulation could be used.

ALGORITHM 5: Quasi-truthful Nash for the Remove mechanism

Input: Voter x , candidates y, z **Output:** One of $\{y, z\}$ **if** x *thinks he should win* **then**

- if** y *would prefer a win for x rather than a three-way tie in a truthful world* **then**
 - return** y ;
- else**
 - return** z ;

else

- return** a truthful comparison between y and z ;

ALGORITHM 6: Quasi-truthful Nash for the RepeatThenRemove mechanism

Input: Voter x , candidates y, z **Output:** One of $\{y, z\}$ **if** x *thinks he should win* **then**

- if** *it is the first TriadicVote* **then**
 - /* For example, a truthful comparison */*
 - return** either of y or z ;
- else**
 - return** the candidate that he didn't vote for in the first round;

else

- return** a truthful comparison between y and z ;

have concave utilities, it turns out that there always must be a participant that prefers y to win over a removal of all three balls. This intuition is translated into a rigorous proof in Sect. 4.

The same idea gives us Alg. 6, a quasi-truthful Nash for the RepeatThenRemove mechanism and a more practical strategy to implement. In this strategy, x simply chooses an arbitrary participant to vote for first; if there is a tie, he should then switch his vote. It turns out that the proof structure for this is almost identical to that of Alg. 5. If y happens to vote for x in the first round, then he is safe according to the reasoning of Alg. 5. However, if y happens to vote for z , the only thing z can do is to cause a repeat vote, during which y will vote for x , bringing us back to the realm of Alg. 5's strategy.

3 Triadic Consensus approximates the Condorcet winner with low communication complexity

3.1 Background: Fixed size urns and urn functions

The primary idea in proving the results in this section is to reduce the Triadic Consensus urn to previously known results for fixed size urns with urn functions. A fixed size urn contains some number of balls, which are each colored either red

or blue. Let R_t and B_t be the number of red and blue balls respectively at time t , where $R_t + B_t = n$. Also, let $p_t = \frac{R_t}{n}$ denote the fraction of red balls. At every discrete time t , either a red ball is sampled with probability $f(p_t)$, a blue ball is sampled with probability $f(1 - p_t)$, or nothing happens with the remaining probability. The function $f : [0, 1] \rightarrow [0, 1]$ is called an urn function and satisfies $0 \leq f(x) + f(1 - x) \leq 1$ for $0 \leq x \leq 1$. If a ball was sampled, it is then recolored to the opposite color and placed back into the urn. This process repeats until some time T when all the balls are the same color, i.e. $R_T = n$ or $R_T = 0$.

We will show in the following section that Triadic Consensus is closely related to fixed size urns with urn function $f(p) = 3p(1 - p)^2$. We will then use the following theorems derived from those in Lee and Bruck [20]⁸.

Theorem 1. *Let a fixed size urn start with R_0 red balls out of n total balls and have an urn function $f(p) = 3p(1 - p)^2$. Let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\Pr[R_T = n] = \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{R_0} \binom{n-1}{j-1}$$

Theorem 2. *Let a fixed size urn start with R_0 red balls out of n total balls and have an urn function $f(p) = 3p(1 - p)^2$. Let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbb{E}[T] \leq n \ln n + O(n)$$

3.2 Reduction from Triadic Consensus to fixed size urns

Recall that our results are for the case of single-peaked preferences, for which the candidates can be said to lie on some axis. Every voter's utility is described by a peak on that axis which falls off on either direction. Without loss of generality, we let the participant ids be labeled from one end of the axis to the other, i.e. $1 < 2 < \dots < n$.

Lemma 1. *Let x , y , and z be three unique participants whose peaks lie on an axis such that $x < y < z$. Then the winner of a quasi-truthful $\text{TriadicVote}(x, y, z)$ must be the median participant y .*

Proof. Since y would win in a truthful vote, this follows from the definition of quasi-truthfulness. \square

Lemma 2. *For single-peaked preferences, let the participant ids be labeled from one end of the axis to the other, i.e. $1 < 2 < \dots < n$. Color balls with ids $1, 2, \dots, i$ red and balls with ids $i + 1, i + 2, \dots, n$ blue. Then if participants vote quasi-truthfully, Triadic Consensus (for $k = 1$) will produce a red winner with the same probability as that of a fixed size urn with urn function $f(p) = 3p(1 - p)^2$.*

⁸ Theorem 2 requires some algebra that may not be immediately clear from the general theorem stated in the reference. For the convenience of the reader, we include these calculations in Appendix B.

Proof. Let p_r and p_b denote the fraction of red and blue balls respectively. Each time three balls are sampled, the median ball must win by Lemma 1, which implies that the majority color must win. Then we have the following four cases:

Three red With probability p_r^3 , there is no change in colors.

Two red, one blue With probability $3p_r^2p_b$, one blue ball is recolored red.

One red, two blue With probability $3p_rp_b^2$, one red ball is recolored blue.

Three blue With probability p_b^3 , there is no change in colors.

These are the transition probabilities for a fixed size urn with i red balls, $n - i$ blue balls, and urn function $f(p) = 3p(1 - p)^2$. Since every transition probability is identical, the final probability of a red winner must be identical. \square

3.3 Main results

Theorem 3. *For single-peaked preferences, let the participant ids be labeled from one end of the axis to the other, i.e. $1 < 2 < \dots < n$. Then if participants vote quasi-truthfully, Triadic Consensus (for $k = 1$) will produce a winner w with probability*

$$\Pr[w = i] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

Proof. If balls $1, 2, \dots, i$ are colored red, then $w \leq i$ iff the winning ball is red. Then applying Theorem 1 and Lemma 2, we get $\Pr[w \leq i]$. By subtracting $\Pr[w \leq i - 1]$ from $\Pr[w \leq i]$, we get our final expression. \square

A similar argument extends the above theorem for general k . Using standard probabilistic arguments[21], we get the following corollary.

Corollary 1. *Let there be n single-peaked participants and let w denote the winning id after running Triadic Consensus with $k = O(\frac{1}{\epsilon^2} \log \frac{1}{\delta})$. Then assuming that participants vote quasi-truthfully, w will be a $(1 - \epsilon/\sqrt{n})$ approximation of the Condorcet winner with probability at least $1 - \delta$.*

Theorem 4. *For single-peaked preferences and quasi-truthful voting, Triadic Consensus has a total communication complexity of $O(kn \log^2(kn))$.*

Proof. Theorem 2 is an upper bound on the expected time it takes to halve the number of remaining participants (since we can color half the participants red and half blue). For kn balls, this gives us less than or equal to $kn \ln(kn) + O(kn)$ expected time to halve the participants. Each time the urn converges to a single color, we can recolor half the remaining participants and repeat. After $\log n$ rounds, we will be done. \square

The above theorem is an upper bound on the communication complexity. In reality, at each recoloring, the balls will not be split evenly between the two colors. Based on this intuition and simulations, we conjecture that the communication complexity is only $O(kn \log kn)$.

Conjecture 1. For single-peaked preferences and quasi-truthful voting, Triadic Consensus has a total communication complexity of $O(kn \log kn)$.

3.4 How bad is $O(\sqrt{n})$ error?

We would like to point out that for crowdsourcing applications, participants can often be viewed as noisy samples from some underlying distribution. Because of this, even if we were to find the exact Condorcet winner of the noisy sample, this may not significantly improve the variance.

For example, suppose that the participants are drawn independently and uniformly from $[0, 1]$ so that the true Condorcet candidate would be one in position $\frac{1}{2}$. Now suppose that $\frac{1}{2}$ lies between the k -th and $k + 1$ -th sampled participant. Then k is clearly binomially distributed, which is the same distribution as the case of Triadic Consensus. In other words, the exact Condorcet winner of the sampled participants has a standard deviation of $\frac{1}{2}\sqrt{n}$ participants between him and the true Condorcet winner of the underlying distribution. Since the approximate Condorcet winner produced by Triadic Consensus also $O(\sqrt{n})$ standard deviation, they have the same order of error.

3.5 Triadic Consensus eliminates outliers quickly

Triadic Consensus has the intuition of quickly eliminating outliers since each participant needs to convince two other participants to vote for him in order to win. We will give two observations supporting this intuition. First, we consider two seemingly powerful algorithms that have access to side information that allows them to directly eliminate outliers in various ways. Despite their use of this knowledge, they choose a winner with the exact same probability distribution as that of Triadic Consensus. Second, we compare Triadic Consensus to a non-Triadic random sampling algorithm and show that Triadic Consensus does much better at picking a central candidate for each round.

Two seemingly powerful algorithms Consider algorithms `RemoveRandomExtreme` (Alg. 7) and `ContractExtremes` (Alg. 8). Each of these algorithms assume a set of participants which have single-peaked preferences and have some limited access to the two extremes (the participants with the leftmost and rightmost peaks). In `RemoveRandomExtreme`, one of the two extreme balls is randomly chosen and thrown out. In `ContractExtremes`, a random ball is chosen, and both of the extreme balls are moved to this chosen ball. We show that they choose a winner with the same probability distribution as that of Triadic Consensus.

Theorem 5. *For single-peaked preferences, let the participant ids be labeled from one end of the axis to the other, i.e. $1 < 2 < \dots < n$. Then `RemoveRandomExtreme` and `ContractExtremes` will produce a winner w with probability*

$$\Pr[w = i] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

Proof. We use the same coloring technique as that of the proof for Lemma 2

ALGORITHM 7: RemoveRandomExtreme

Input: An urn with a ball for each participant $1, 2, \dots, n$, each of whom lie on an axis

Output: A winning candidate i .

while *there is more than one label* **do**

 Randomly sample one of the two (left or right) extreme balls;
 Toss out the sampled ball;

return remaining label;

ALGORITHM 8: ContractExtremes

Input: An urn with a ball for each participant $1, 2, \dots, n$, each of whom lie on an axis

Output: A winning candidate i .

while *there is more than one label* **do**

 Randomly sample one candidate;
 Move both of the two (left and right) extreme balls to the sampled ball;

return remaining label;

Hot-or-Not Consensus Consider Hot-or-Not Consensus, in which two balls are randomly chosen as candidates and one single ball is randomly chosen as the voter. The voter then votes between the two chosen balls and the two candidate balls are replaced with the winning candidate. We do a one-step comparison of these two algorithms given a continuous distribution of participants. It turns out that the naive Hot-or-Not Consensus can be thought of as a mix between Triadic Consensus and the (really bad) method of randomly picking a candidate.⁹

Theorem 6. *Let a continuous distribution of voters be uniformly distributed between zero and one. Let $g_{\text{Hot-or-Not}}(x)$ and $g_{\text{Triadic}}(x)$ be the probability density of x being the next winning candidate in Hot-or-Not and Triadic Consensus respectively. Then, $g_{\text{Triadic}} = 6x(1-x)$ and $g_{\text{Hot-or-Not}} = 3x(1-x) + \frac{1}{2}$. In particular,*

$$g_{\text{Hot-or-Not}} = \frac{1}{2}g_{\text{Triadic}} + \frac{1}{2}g_{\text{Unif}}$$

where g_{Unif} is the uniform distribution over the interval $[0, 1]$.

Proof. For uniformly distributed participants, we have the density function $f(x) = 1$ and cumulative density function $F(x) = x$. In Triadic Consensus, x wins if he is selected along with a candidate to the left and right of him. Then, we have

$$g_{\text{Triadic}}(x) = 3!f(x)F(x)(1-F(x)) = 6x(1-x)$$

⁹ This is only an intuition based on a one-step comparison and should not be interpreted as a comparison of their final approximation values.

In Hot-or-Not Consensus, x wins against y only if the voter z is closer to x than y . Then,

$$\begin{aligned} g_{\text{Hot-or-Not}}(x) &= 2f(x) \int_0^x f(y) \left(1 - F\left(\frac{x+y}{2}\right)\right) dy + 2f(x) \int_x^1 f(y) F\left(\frac{x+y}{2}\right) dy \\ &= 2 \left[\int_0^x \left(1 - \frac{x+y}{2}\right) dx + \int_x^1 \left(\frac{x+y}{2}\right) dx \right] = 3x(1-x) + \frac{1}{2} \end{aligned}$$

3.6 Simulations in general spaces show promise

Even though our approximation and communication complexity results only hold for single-peaked preferences, we believe that Triadic Consensus has strong properties for more complex spaces. To demonstrate this, we consider two classes of preferences that are induced by points in a two dimensional Euclidean space. In each of these cases, the rankings are clearly not single-peaked, but simulations show strong results for both the approximation and the communication complexity. The first example is a straightforward generalization to points that are laid out in a grid.

Example 3. $n = m^2$ voters are placed on the points $(0, 0), (0, 1), \dots, (m-1, m-1)$ to form a $m \times m$ grid. The Condorcet winner in this scenario is at the median point $(\frac{m}{2}, \frac{m}{2})$. From the simulation results below, we can see that Triadic Consensus picks winners that are closely distributed around the winner. The average number of votes each voter casts is $\sim O(\log n)$.

	mean winner	σ of winner	mean votes/voter	σ of votes/voter
5x5	(1.96, 1.97)	.953	1.966	0.425
10x10	(4.45, 4.5)	1.157	3.405	0.517
20x20	(9.61, 9.47)	1.236	4.618	0.417
40x40	(19.64, 19.53)	1.594	6.056	0.368
80x80	(39.57, 39.69)	1.555	7.293	0.324

Table 1: Simulation results for 100 iterations of Triadic Consensus on a grid

In the second case, we try to design a difficult scenario by densely populating the perimeter of a circle and adding a single point at its center, who is the Condorcet winner.

Example 4. n voters are placed uniformly around a circle (in the plane) with radius 1 and centered at $(0, 0)$. A single voter is placed at the point $(0, 0)$. The Condorcet winner in this scenario is the point $(0, 0)$. Surprisingly, we find that even as the number of points increases on the perimeter, the probability of the randomized algorithm selecting the single Condorcet winner still remains non-trivial. If this probability does remain above some constant, then we can

use standard probabilistic methods to show that repeating Triadic Consensus a small number of times and picking the majority winner will result in $(0,0)$ with arbitrarily high accuracy. The average number of votes each voter casts is $\sim O(\log n)$.

	% times $(0,0)$ wins	mean votes/voter	σ of votes/voter
25	0.368	2.225	0.513
100	0.338	4.189	0.775
400	0.305	6.628	1.185
1600	0.295	9.224	1.696
6400	0.302	11.764	2.321

Table 2: Simulation results for 1000 iterations of Triadic Consensus on a circle with a single point in the center

4 Triadic Consensus has a quasi-truthful Nash equilibrium for concave utilities

To discuss strategic behavior, we need to define the utilities for each participant. Let $U_x(y)$ denote the utility that x gets from a proposal y . The utility x derives from y depends on the distance from x to y , i.e. $U_x(y) = f_x(d(x, y))$, where $f(\cdot)$ must be decreasing in distance so that $U_x(y) > U_x(z)$ whenever $y \succ_x z$. We say that a participant x has a concave utility function if $f_x(\cdot)$ is a concave function.

Theorem 7. *If all participants have concave utility functions, then Algs. 5 and 6 are quasi-truthful Nash equilibria for Triadic Consensus when using the Remove and RepeatThenRemove mechanisms respectively.*

Proof. We prove our main result with the following proof by induction. Since the proofs for the Remove mechanism and the RepeatThenRemove mechanism are almost identical (see 2.2), we will refer solely to the Remove mechanism for simplicity.

Base Case: Alg. 5 is a Nash equilibrium for $n = 1, 2, 3$ balls (Lemma 3).

Inductive Step: Assume that Alg. 5 is a Nash equilibrium for $n - 3$ balls. Now consider a participant x who is considering deviating from Alg. 5 in an urn with n balls:

1. For any TriadicVote with participants $x < y < z$ in an urn with n balls, if y votes for x , then by the definition of the strategy and the fact that one of x and z must prefer y to a three-way tie (Lemma 4), we know that x must prefer y to win over a three-way tie, which means x should not deviate.

2. For any TriadicVote with participants $x < y < z$ in an urn with n balls, if y votes for z , then given that the previous statement is true, we show that x should prefer a win for y over a win for z (Lemma 5). This is done by defining a comparison relation between urns that formalizes this intuition that participants should prefer closer balls. With this definition, we can define a coupling of two urns: one in which x plays an optimal strategy, and one in which x always plays according to Alg. 5. We show that for every coupled history, the urn from Alg. 5 does at least as well as the optimal urn in expected utility. This means that Alg. 5 is also an optimal strategy for x in this case.

By carrying out the Inductive Hypothesis, we get our result for all n . \square

4.1 Supporting Lemmas

Lemma 3. *Alg. 5 is a Nash equilibrium for Triadic Consensus with the Remove mechanism when $n = 1, 2$, or 3 balls.*

Proof. This is trivially true for $n = 1$ and 2 since no votes take place. For $n = 3$, suppose that the three participants are $x < y < z$. In this case, the only situation when participants cast votes is when TriadicVote is performed with all three unique participants. After such a situation occurs, there will either be a winner or all balls will be eliminated and no further votes take place. Therefore, our analysis can be constrained to this single TriadicVote.

If participants vote according to Alg. 5, we know that y will be the winner since x and z both vote for him. Suppose y votes for x WLOG. Then if z deviates, x will win, which is clearly suboptimal. If x deviates, then there is a three-way tie and all are eliminated, resulting in a uniformly random winner.

The difference in utility lost for x by deviating is $\Delta U_x = U_x(y) - \frac{1}{3}(U_x(x) + U_x(y) + U_x(z))$. Letting d_1 be the distance between x and y and d_2 the distance between y and z , we have $\Delta U_x = \frac{1}{3}(f_x(d_1) - f_x(0)) - \frac{1}{3}(f_x(d_1 + d_2) - f_x(d_1))$ and

$$\begin{aligned} \Delta U_x \geq 0 &\iff f_x(d_1) - f_x(0) \geq f_x(d_1 + d_2) - f_x(d_1) \\ &\iff \frac{f_x(d_1) - f_x(0)}{f_x(d_1 + d_2) - f_x(d_1)} \leq 1 \end{aligned}$$

Similarly,

$$\Delta U_z \geq 0 \iff \frac{f_z(d_2) - f_z(0)}{f_z(d_1 + d_2) - f_z(d_2)} \leq 1$$

For concave, monotonically non-increasing f_x and f_z , we know that (detailed in the long version[?]):

$$\frac{f_x(d_1) - f_x(0)}{f_x(d_1 + d_2) - f_x(d_1)} \leq \frac{d_1}{d_2} \quad \text{and} \quad \frac{f_z(d_2) - f_z(0)}{f_z(d_1 + d_2) - f_z(d_2)} \leq \frac{d_2}{d_1}$$

But then, at least one of $\frac{d_1}{d_2}$ or $\frac{d_2}{d_1}$ is less than or equal to 1, which means that at least one of ΔU_x and ΔU_z is greater than or equal to 0 and prefers a win for y over a three-way tie. By the definition of Alg. 5, y will vote for this person when he exists. Therefore, since y voted for x , we know $\Delta U_x \geq 0$, which concludes the proof. \square

Lemma 4. *Assume that Alg. 5 is a Nash equilibrium for any configuration of $n-3$ balls. Then for a TriadicVote among participants $x < y < z$ in an urn with n balls, at least one of x or z prefers a win for y over a three-way tie, so long as they both have concave utilities.*

Proof. Because of space constraints, we will only outline the proof here, leaving the notation and algebra for Appendix A. The proof has two parts:

Part A. Suppose all balls are positioned somewhere between x and z , i.e. in the interval $[x, z]$. Then, if x and z have concave utility functions, at least one of x and z prefers a win for y over a three-way tie. The proof for this statement is similar to the one in Lemma 3, albeit more complex.

Part B. For any configuration of n balls, moving any ball at position x leftwards and moving any ball at position z rightwards can only increase both ΔU_x and ΔU_z . Put another way, given any configuration, we can move all balls left of x to x and all balls right of z to z , while only decreasing ΔU_x and ΔU_z . Once moved in this way, the configuration of balls falls under the jurisdiction of Part 1, which states that at least one of ΔU_x and ΔU_z is greater than or equal to 0. Therefore, the same participant in the original configuration must also have a positive ΔU , which means he prefers a win for y over a three-way tie. \square

For the final lemma, we require the following definition.

Definition 1. *Given two urns R and S , each with n balls, number the balls in R from left to right as r_1, r_2, \dots, r_n and number the balls in S from left to right as l_1, l_2, \dots, l_n . Then R x -dominates S if*

$$s_i \leq r_i \text{ for } r_i < x$$

$$s_i = r_i \text{ for } r_i = x$$

$$s_i \geq r_i \text{ for } r_i > x$$

Lemma 5. *Assume that Alg. 5 is a Nash equilibrium for any configuration of $n-3$ balls. Then for a TriadicVote among participants $x < y < z$ in an urn with n balls, if y votes for z (WLOG), x does not benefit by voting strategically for z .*

Proof. Our proof strategy will be to use a coupling argument. Let OPT denote the optimal strategy for x . We consider two urns R and S . In urn R , x plays according to Alg. 5. In urn S , x plays according to OPT, the strategy that maximizes his expected utility. We couple the TriadicVote's of these urns in the following way:

1. Let r_1, r_2, \dots, r_n denote the balls in urn R as indexed from leftmost position to rightmost position. Let s_1, s_2, \dots, s_n denote the balls in urn S as indexed from leftmost position to rightmost position.

2. Then for every TriadicVote, when balls r_i, r_j, r_k are randomly drawn from urn R , balls s_i, s_j, s_k will be drawn from urn S .

Suppose R x -dominates S and then each undergoes a coupled TriadicVote where balls $r_i < r_j < r_k$ are selected from R and $s_i < s_j < s_k$ are selected from S . After they vote, we show that the resulting urns R' and S' must still satisfy R' x -dominates S' . By the coupling rule, this is trivially true when 1) x is not selected, 2) x is represented in two or more balls, and 3) x is the middle participant. This is because x either does not vote or cannot affect the result in these cases (remember that all other participants are voting according to Alg. 5). The only remaining case is when x is one of the side participants (s_i WLOG). In this case, r_j wins in urn R since x plays according to Alg. 5 in this urn. Suppose s_j voted for s_k . Then regardless of who x votes for, one of s_j or s_k must win, both of which will still satisfy R' x -dominates S' . Now suppose s_j voted for s_i . Then x could eliminate all three participants by voting for s_k . However, by Lemma 4 and the definition of Alg. 5, this would be suboptimal, which means that x cannot play this strategy in urn S . Therefore, R' x -dominates S' .

Finally, we note that before any TriadicVote's take place, R and S are identical, i.e. R x -dominates S . Then, the winner of R must also x -dominate the winner of S , which means that urn R is better for x in every coupled history. \square

5 Future Directions

There are many future directions for this work.

Triadic Consensus For the algorithm itself, the primary problem that begs to be worked on is an analysis for higher dimensional or even non-Euclidean spaces. It is an open question whether Triadic Consensus achieves low communication complexity for general preference profiles and whether it has nice properties of convergence. The authors believe that there is something interesting that can be said here. Similarly, it would also be exciting to extend the work on quasi-truthfulness to higher dimensional spaces. The authors do not believe that a naive extension will suffice; however, it seems possible that probabilistic strategies coupled with other punishments for manipulation will be able to achieve this goal.

Truthful voting rules When participants are voters and candidates, we have indicated that manipulation can often be detected. It would be interesting to use this idea, possibly along with the theme of triads, quasi-truthfulness, and cyclic preferences, to design truthful voting rules. For example, one could imagine the following variant of the Borda count: for each of the $\binom{n}{3}$ triads, add one point to the score of the winner¹⁰.

¹⁰ The Borda count is equivalent to giving the winner two points, the next highest scoring participant one point, and the loser zero points.

Communication complexity Another exciting problem is to make new approximate and randomized voting rules that have low communication complexity. In particular, it would be useful to have a voting rule where the maximum number of comparisons per voter is small (say, $O(\log n)$). In Triadic Consensus, only the average number of comparisons is small, which may still prevent it from being widely applicable to large internet crowdsourcing applications.

Consensus mechanisms On the direction of group consensus mechanisms, one possible extension of this work is to bring it outside of voting. Namely, rather than having the randomly selected triads *vote*, it would be interesting to analyze other sorts of dynamics that are more collaborative or game-theoretic.

Urn voting rules It would also be interesting to study generalized urn voting rules. This could include different ball replacement schemes or even more elaborate generalizations. For example, balls could be labeled with participant *and* proposal ids so that only proposal ids are changed after a TriadicVote. Such urn voting rules are interesting because they can be interpreted as local decisions made by small groups of people and also provide a natural framework for studying (non-trivial) probabilistic voting rules.

Acknowledgements

The authors would like to thank Pranav Dandekar for helpful discussions and Vincent Conitzer for pointing us to several pieces of related literature and posing the idea of lowering communication complexity with approximations. This research was supported in part by NSF grants 0904325 and 0947670. David Lee was also supported in part by an NSF Graduate Research Fellowship.

References

1. Brandt, F., Conitzer, V., Endriss, U.: Computational social choice. In: Multiagent Systems. MIT Press (2012)
2. Young, H.P.: Condorcet’s theory of voting. *The American Political Science Review* **82**(4) (1988) 1231–1244
3. Arrow, K.: A difficulty in the concept of social welfare. *Journal of Political Economy* **58**(4) (1950) 328–346
4. Pattanaik, P., Peleg, B.: Distribution of power under stochastic social choice rules. *Econometrica* **54**(4) (1986) 909–921
5. Gibbard, A.: Manipulation of voting schemes: A general result. *Econometrica* **41**(4) (1973) 587–601
6. Satterthwaite, M.: Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* **10** (1975) 187–217
7. Gibbard, A.: Manipulation of schemes that mix voting with chance. *Econometrica* **45**(3) (1977) 665–681

8. Bartholdi, J., Tovey, C., Trick, M.: The computational difficulty of manipulating an election. *Social Choice and Welfare* **6** (1989) 227–241
9. Faliszewski, P., Procaccia, A.: Ai’s war on manipulation: Are we winning? *AI Magazine* **31**(4) (2010) 53–64
10. Conitzer, V., Sandholm, T.: Nonexistence of voting rules that are usually hard to manipulate. In: *Proceedings of the 21st AAAI Conference*. (2006) 627–634
11. Procaccia, A.: Can approximation circumvent gibbard-satterthwaite? In: *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. (2010) 836–841
12. Birrell, E., Pass, R.: Approximately strategy-proof voting. In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. (2011)
13. Conitzer, V., Sandholm, T.: Communication complexity of common voting rules. In: *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*. (2005) 78–87
14. Conitzer, V., Sandholm, T.: Vote elicitation: Complexity and strategy-proofness. In: *Proceedings of the 17th AAAI Conference*. (2002)
15. Lu, T., Boutilier, C.: Robust approximation and incremental elicitation in voting protocols. In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. (2011)
16. Black, D.: On the rationale of group decision-making. *Journal of Political Economy* **56**(1) (1948) 23–34
17. Moulin, H.: On strategy-proofness and single peakedness. *Public Choice* **35**(4) (1980) 437–455
18. Escoffier, B., Lang, J., Öztürk, M.: Single-peaked consistency and its complexity. In: *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*. (2008)
19. Conitzer, V.: Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research* **35** (2009) 161–191
20. Lee, D., Bruck, J.: Modeling biological circuits with urn functions. In: *Proceedings of the International Symposium on Information Theory (ISIT)*. (2012)
21. Motwani, R., Raghavan, P.: *Randomized Algorithms*. Cambridge University Press (1995)

A Proof of Lemma 4

We restate Lemma 4 here for the convenience of the reader.

Lemma 6. *Assume that Alg. 5 is a Nash equilibrium for any configuration of $n - 3$ balls. Then for a TriadicVote among participants $x < y < z$ in an urn with n balls, at least one of x or z prefers a win for y over a three-way tie, so long as they both have concave utilities.*

As described in Section 4, this proof will be broken up into two parts, which we will prove below. Each of these lemmas assume the inductive hypothesis that Alg. 5 is a Nash equilibrium for any configuration of $n - 3$ balls.

A.1 Preliminaries

Let P denote the urn resulting from a win for y and Q denote the urn resulting from a three-way tie. P has n balls, while Q has $n - 3$ balls.

Index the balls in urn P as b_1, b_2, \dots, b_n from the leftmost participant position to the rightmost and let the balls in urn Q be indexed identically. Let b_l, b_{l+1}, b_{l+2} denote the three balls in urn P that are not in urn Q , so that urn Q has balls $b_1, \dots, b_{l-1}, b_{l+3}, \dots, b_n$. Note that these three balls can be indexed consecutively because all of these missing balls are at the same location (participant y).

For each ball b_i , there is some probability that the ball wins in urn P and some probability that the ball wins in urn Q .¹¹ Since b_l, b_{l+1}, b_{l+2} don't exist in urn Q , their probability of winning there is simply 0.

Let $\Delta p(b_i) = \Pr[b_i \text{ wins in urn } P] - \Pr[b_i \text{ wins in urn } Q]$.

We will use $U_x(b_i)$ to denote the utility of a b_i win for x . As before, $U_x(b_i) = f_x(d(x, b_i))$. If f_x is concave, then x is said to have a concave utility function. Let U_x^P denote the expected utility for x from quasi-truthful voting in urn P . Finally, we use $\Delta U_x = U_x^P - U_x^Q$ to denote the difference in expected utility in quasi-truthful voting for urn P and compared to urn Q .

With this notation, our proof essentially boils down to proving that at least one of $\Delta U_x \geq 0$ or $\Delta U_z \geq 0$ holds. We start with a Lemma which we will need for the further arguments. Roughly speaking, it states that there is an interval from a ball on participant y to the median ball for which the probability of winning in P is greater than Q . All balls outside this interval have a lower probability of winning in P than in Q .

Lemma 7. *Suppose Triadic Consensus is run on the urns P and Q . Then for quasi-truthful voting,*

$$\begin{aligned} \Delta p(b_i) &> 0 \text{ if } \min(l, n/2) \leq i \leq \max(l+2, n/2) \\ \Delta p(b_i) &< 0 \text{ otherwise} \end{aligned}$$

Proof. Recall that $\Delta p(b_i) = \Pr[b_i \text{ wins in urn } P] - \Pr[b_i \text{ wins in urn } Q]$. From Theorem 3, we have that a quasi-truthful strategy in urns P and Q give,

$$\begin{aligned} \Pr[b_i \text{ wins in urn } P] &= \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1} \\ \Pr[b_i \text{ wins in urn } Q] &= \begin{cases} \left(\frac{1}{2}\right)^{n-4} \binom{n-4}{i-1} & \text{if } i \leq l-1 \\ \left(\frac{1}{2}\right)^{n-4} \binom{n-4}{i-4} & \text{if } i \geq l+3 \\ 0 & \text{if } i = l, l+1, l+2 \end{cases} \end{aligned}$$

Note that for $i < l$,

$$\begin{aligned} \Delta p(b_i) &= \left(\frac{1}{2}\right)^{n-1} \left[\frac{(n-1)!}{(i-1)!(n-i)!} - 8 \frac{(n-4)!}{(i-1)!(n-i-3)!} \right] \\ &= \left(\frac{1}{2}\right)^{n-1} \frac{(n-4)!}{(i-1)!(n-i)!} \underbrace{[(n-1)(n-2)(n-3) - 8(n-i)(n-i-1)(n-i-2)]}_{f(i)} \end{aligned}$$

¹¹ Technically, it is not clear what it means for one ball in position p to win over another ball in position p ; we let the winning probability of the i -th ball of n total balls be $\propto \binom{n-1}{i-1}$ since this is convenient and still produces the correct participant winning probabilities.

Since $f(i)$ is monotonically increasing in i , then by observing that $f(\frac{n}{2} - 1) < 0$ and $f(\frac{n}{2}) > 0$, we have

$$\Delta p(b_i) \text{ is } \begin{cases} < 0 & \text{if } i < \min(l, n/2) \\ > 0 & \text{if } n/2 \leq i < l \end{cases}$$

Similarly, for $i > l + 2$, we can use an analogous argument (or apply symmetry) to claim that,

$$\Delta p(b_i) \text{ is } \begin{cases} < 0 & \text{if } i > \max(l + 2, n/2) \\ > 0 & \text{if } l + 2 < i \leq n/2 \end{cases}$$

Finally, it is clear that $\Delta p(b_i) > 0$ for $i = l, l + 1, l + 2$, so we are done.

A.2 Part A: If all b_i lie in $[x, z]$

Lemma 8. *Assume that the inductive hypothesis holds for $n - 3$. Then for a TriadicVote among participants $x < y < z$ in an urn with n balls, each of which lie in $[x, z]$, at least one of x or z prefers a win for y over a three-way tie, so long as they both have concave utilities.*

Recall that our lemma boils down to proving that at least one of $\Delta U_x \geq 0$ and $\Delta U_z \geq 0$ is true. Since the proof is very notation heavy, we first sketch the proof with an example.

Example 5. Let P be an urn with five balls: b_1 is located at position x ; b_2, b_3 , and b_4 are located at position y ; and b_5 is located at position z . Then urn Q is an urn with the two balls b_1 and b_5 .

We know (see Theorem 3) that for balls b_1, b_2, b_3, b_4 , and b_5 , $\Pr[b_i \text{ wins in urn } P]$ is $\frac{1}{16}, \frac{4}{16}, \frac{6}{16}, \frac{4}{16}$, and $\frac{1}{16}$ respectively; $\Pr[b_i \text{ wins in urn } Q]$ is $\frac{1}{2}, 0, 0, 0$, and $\frac{1}{2}$, respectively; which means that $\Delta p(b_i)$ is $-\frac{7}{16}, \frac{4}{16}, \frac{6}{16}, \frac{4}{16}$, and $-\frac{7}{16}$, respectively. Then,

$$\begin{aligned} \Delta U_x &= \Delta p(b_1)U_x(b_1) + \left[\sum_{i=2}^4 \Delta p(b_i)U_x(b_i) \right] + \Delta p(b_5)U_x(b_5) \\ &= -\frac{7}{16}f_x(d(x, x)) + \frac{14}{16}f_x(d(x, y)) - \frac{7}{16}f_x(d(x, z)) \\ &= -\frac{7}{16}[f_x(d(x, x)) - f_x(d(x, y))] - \frac{7}{16}[f_x(d(x, z)) - f_x(d(x, y))] \end{aligned}$$

Since $\frac{7}{16}[f_x(d(x, x)) - f_x(d(x, y))] \geq 0$, we have that

$$\Delta U_x \geq 0 \iff \frac{-\frac{7}{16}[f_x(d(x, z)) - f_x(d(x, y))]}{-\frac{7}{16}[f_x(d(x, x)) - f_x(d(x, y))]} \geq 1$$

Using similar arguments and the fact that $\frac{7}{16}[f_z(d(z, x)) - f_z(d(z, y))] \leq 0$, we have that

$$\Delta U_z \geq 0 \iff \frac{-\frac{7}{16}[f_z(d(z, z)) - f_z(d(z, y))]}{-\frac{7}{16}[f_z(d(z, x)) - f_z(d(z, y))]} \leq 1$$

By concavity and Lemma 10 in Appendix B,

$$\begin{aligned} \frac{f_x(d(x, z)) - f_x(d(x, y))}{f_x(d(x, x)) - f_x(d(x, y))} &\geq \frac{(z - x) - (y - x)}{(x - x) - (y - x)} = \frac{z - y}{x - y} \\ &= \frac{(z - z) - (z - y)}{(z - x) - (z - y)} \geq \frac{f_z(d(z, z)) - f_z(d(z, y))}{f_z(d(z, x)) - f_z(d(z, y))} \end{aligned}$$

This means that if $\frac{z-y}{x-y} \geq 1$, then $\Delta U_x \geq 0$. Otherwise, if $\frac{z-y}{x-y} \leq 1$, then $\Delta U_z \geq 0$. Obviously, one of these must be true, so we are done for this example.

Proof. We will now prove our lemma for the general case. The argument structure is exactly the same. Let $A = \min(l, n/2)$ and $Z = \max(l + 2, n/2)$. Recall that these are the leftmost and rightmost balls for which $\Delta p(b_i) > 0$. All other balls must have $\Delta p(b_i) < 0$. Then we can separate the expression for ΔU_x into three summations,

$$\Delta U_x = \sum_{i=1}^{A-1} \Delta p(b_i) U_x(b_i) + \sum_{i=A}^Z \Delta p(b_i) U_x(b_i) + \sum_{i=Z+1}^n \Delta p(b_i) U_x(b_i)$$

Note that $\sum_{i=1}^{A-1} \Delta p(b_i) + \sum_{i=Z+1}^n \Delta p(b_i) = -\sum_{i=A}^Z \Delta p(b_i)$. Then we can partition up the mass of $\Delta p(b_j)$ so that we can get one-to-one correspondence of masses corresponding to $j \in [A, Z]$ and $j \notin [A, Z]$. Let the masses corresponding to $i = 1, 2, \dots, A-1$ be denoted by p_1, p_2, \dots, p_u and let the masses corresponding to Z_1, Z_2, \dots, n be denoted by q_1, q_2, \dots, q_v . In other words, $p_{1..u}, q_{1..v}$ are chosen so that $\sum_{i=1}^{A-1} \Delta p(b_i) = \sum_{i=1}^u p_i$, $\sum_{i=Z+1}^n \Delta p(b_i) = \sum_{i=1}^v q_i$, and

$$\Delta U_x = \sum_{i=1}^u p_i [U_x(b_{g_1(i)}) - U_x(b_{h_1(i)})] + \sum_{i=1}^v q_i [U_x(b_{g_2(i)}) - U_x(b_{h_2(i)})]$$

where $g_1 : [1..u] \rightarrow [1..A-1]$, $g_2 : [1..v] \rightarrow [Z+1..n]$, $h_1 : [1..u] \rightarrow [A..Z]$, and $h_2 : [1..v] \rightarrow [A..Z]$.

Note that since balls are indexed left to right and all lie within $[x, z]$, then balls indexed $i \in [1, A)$ (e.g. $b_{g_1(i)}$) are closer to x than those indexed $i \in [A, Z]$ (e.g. $b_{h_1(i)}$ and $b_{h_2(i)}$), which are closer than those indexed $i \in (Z, n]$ (e.g. $b_{g_2(i)}$). Therefore, we have

$$U_x(b_{g_1(i)}) - U_x(b_{h_1(i)}) \geq 0 \text{ and } U_x(b_{g_2(i)}) - U_x(b_{h_2(i)}) \leq 0$$

Combining these, we get that

$$\Delta U_x \geq 0 \iff \frac{\sum_{i=1}^u p_i [U_x(b_{g_1(i)}) - U_x(b_{h_1(i)})]}{\sum_{i=1}^v q_i [U_x(b_{h_2(i)}) - U_x(b_{g_2(i)})]} \geq 1 \quad (1)$$

since $q_i[U_x(b_{h_2(i)}) - U_x(b_{g_2(i)})] \geq 0$. Similarly, we have

$$\Delta U_z = \sum_{i=1}^u p_i[U_z(b_{g_1(i)}) - U_z(b_{h_1(i)})] + \sum_{i=1}^v q_i[U_z(b_{g_2(i)}) - U_z(b_{h_2(i)})]$$

For z , balls indexed $i \in [1, A)$ (e.g. $b_{g_1(\cdot)}$) are farther from z than those indexed $i \in [A, Z]$ (e.g. $b_{h_1(\cdot)}$ and $b_{h_2(\cdot)}$), which are farther than those indexed $i \in (R, n]$ (e.g. $b_{g_2(\cdot)}$). Therefore,

$$U_z(b_{g_1(i)}) - U_z(b_{h_1(i)}) \leq 0 \text{ and } U_z(b_{g_2(i)}) - U_z(b_{h_2(i)}) \geq 0$$

Combining these, we get that

$$\Delta U_z \geq 0 \iff \frac{\sum_{i=1}^u p_i[U_z(b_{g_1(i)}) - U_z(b_{h_1(i)})]}{\sum_{i=1}^v q_i[U_z(b_{h_2(i)}) - U_z(b_{g_2(i)})]} \leq 1 \quad (2)$$

since $q_i[U_z(b_{h_2(i)}) - U_z(b_{g_2(i)})] \leq 0$.

We now have one last step. For any $f(x)$ which is concave and monotonically decreasing, we have that

$$\frac{\sum_{i=1}^m c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^n d_j[f(s_j^2) - f(s_j^1)]} \geq \frac{\sum_{i=1}^m c_i[t_i^2 - t_i^1]}{\sum_{j=1}^n d_j[s_j^2 - s_j^1]} \quad \text{and} \quad \frac{\sum_{i=1}^m c_i[f(s_i^1) - f(s_i^2)]}{\sum_{j=1}^n d_j[f(t_j^1) - f(t_j^2)]} \leq \frac{\sum_{i=1}^m c_i[s_i^1 - s_i^2]}{\sum_{j=1}^n d_j[t_j^1 - t_j^2]}$$

for $s_j^1 \leq t_i^1$, $s_j^2 \leq t_i^2$, $s_i^1 \leq s_j^2$, $t_i^1 \leq t_j^2$, and $\text{sign}(c_i) = \text{sign}(d_j)$ (as detailed in Appendix B). Applying this to (1) and (2), we get

$$\frac{\sum_{i=1}^u p_i[U_x(b_{g_1(i)}) - U_x(b_{h_1(i)})]}{\sum_{i=1}^v q_i[U_x(b_{h_2(i)}) - U_x(b_{g_2(i)})]} \leq \frac{\sum_{i=1}^u p_i[d(x, b_{g_1(i)}) - d(x, b_{h_1(i)})]}{\sum_{i=1}^v q_i[d(x, b_{h_2(i)}) - d(x, b_{g_2(i)})]} = \frac{\sum_{i=1}^u p_i[-d(b_{g_1(i)}, b_{h_1(i)})]}{\sum_{i=1}^v q_i[-d(b_{h_2(i)}, b_{g_2(i)})]}$$

and

$$\frac{\sum_{i=1}^u p_i[U_z(b_{g_1(i)}) - U_z(b_{h_1(i)})]}{\sum_{i=1}^v q_i[U_z(b_{h_2(i)}) - U_z(b_{g_2(i)})]} \geq \frac{\sum_{i=1}^u p_i[d(z, b_{g_1(i)}) - d(z, b_{h_1(i)})]}{\sum_{i=1}^v q_i[d(z, b_{h_2(i)}) - d(z, b_{g_2(i)})]} = \frac{\sum_{i=1}^u p_i[d(b_{g_1(i)}, b_{h_1(i)})]}{\sum_{i=1}^v q_i[d(b_{h_2(i)}, b_{g_2(i)})]}$$

Then either

$$\frac{\sum_{i=1}^u p_i[d(b_{g_1(i)}, b_{h_1(i)})]}{\sum_{i=1}^v q_i[d(b_{h_2(i)}, b_{g_2(i)})]} \geq 1 \quad \text{or} \quad \frac{\sum_{i=1}^u p_i[d(b_{g_1(i)}, b_{h_1(i)})]}{\sum_{i=1}^v q_i[d(b_{h_2(i)}, b_{g_2(i)})]} \leq 1$$

If the first is true, we can apply (1) to claim that $\Delta U_x \geq 0$. If the second is true, we can apply (2) to claim that $\Delta U_z \geq 0$. Since one of these must be true, we are done.

A.3 Part B: Moving balls into the interval $[x, z]$ only decreases ΔU_x and ΔU_z

Lemma 9. *Consider urns P and Q . From P , create a new urn P' by moving all balls left of x to x and all balls right of z to z . Similarly, from Q , create a new urn Q' by moving all balls left of x to x and all balls right of z to z . Let $\Delta p'(b_i) = \Pr[b_i \text{ wins in urn } P'] - \Pr[b_i \text{ wins in urn } Q']$ and $\Delta U'_x = U_x^{P'} - U_x^{Q'}$. Then,*

$$\Delta U_x \geq \Delta U'_x \quad \text{and} \quad \Delta U_z \geq \Delta U'_z$$

Proof. Since the relative positions of the balls have not changed in P' and Q' , the change in winning probabilities from P' to Q' are the same, i.e. $\Delta p(b_i) = \Delta p'(b_i)$. Then for $\Delta p(b_i) < 0$,

$$\begin{aligned} \Delta p(b_i)U_x(b_i) \geq \Delta p'(b_i)U'_x(b_i) &\iff U_x(b_i) \leq U'_x(b_i) \\ &\iff b_i \text{ is moved closer to } x \end{aligned}$$

and for $\Delta p(b_i) > 0$,

$$\begin{aligned} \Delta p(b_i)U_x(b_i) \geq \Delta p'(b_i)U'_x(b_i) &\iff U_x(b_i) \geq U'_x(b_i) \\ &\iff b_i \text{ is moved farther from } x \end{aligned}$$

Likewise, the same statements hold when x is replaced with z . Then we can use this to prove our Lemma by splitting it up into three cases:

Case 1: $b_{n/2} \in [x, z]$ for urn P When the median ball is located in $[x, z]$, we know from Lemma 7 that all balls left of x and right of z satisfy $\Delta p(b_i) < 0$. Then moving the balls left of x to x brings them closer to both x and z . Similarly, moving the balls right of z to z brings them closer to both x and z . Therefore, we must have

$$\Delta U_x = \sum_{i=1}^n \Delta p(b_i)U_x(b_i) \geq \sum_{i=1}^n \Delta p'(b_i)U'_x(b_i) = \Delta U'_x$$

and similarly, $\Delta U_z \geq \Delta U'_z$.

Case 2: $b_{n/2} < x$ for urn P When the median ball is located to the left of x , we will need to first make an intermediate pair of urns. Note that the balls left of $b_{n/2}$ have $\Delta p(b_i) < 0$ so we can move them rightwards to bring them closer to both x and z . However, since the balls between $b_{n/2}$ and x have $\Delta p(b_i) > 0$, we need to move them leftwards to bring them farther away from both x and z . For the balls right of z , $\Delta p(b_i) < 0$, so we can again move them to z which brings them closer to x and z . We will denote these movements with the intermediate urns P'' and Q'' . In these urns, all the balls left of x are moved to the position $b_{n/2}$ and all the balls right of z are moved to z . By an argument similar to that

of Case 1, $\Delta U_x \geq \Delta U_x''$ and $\Delta U_z \geq \Delta U_z''$. Now, note that

$$\Delta U_x'' = \left[\sum_{i=1}^{k-1} \Delta p''(b_i) \right] U_x''(b_{n/2}) + \sum_{i=k}^n \Delta p''(b_i) U_x''(b_i)$$

where k is the index of the leftmost ball that is right of x and U'', p'' are the analogous expressions for utility and winning probability in urns P'', Q'' . But we also know that for $k < l$,

$$\begin{aligned} \sum_{i=1}^{k-1} \Delta p''(b_i) &= \sum_{i=1}^{k-1} \left[\left(\frac{1}{2} \right)^{n-1} \binom{n-1}{i-1} - \left(\frac{1}{2} \right)^{n-4} \binom{n-4}{i-1} \right] \\ &= \Pr[\leq k-2 \text{ heads in } n-1 \text{ coin flips}] - \Pr[\leq k-2 \text{ heads in } n-4 \text{ coin flips}] \leq 0 \end{aligned}$$

Then this means that

$$\begin{aligned} \left[\sum_{i=1}^{k-1} \Delta p''(b_i) \right] U_x''(b_{n/2}) &\geq \left[\sum_{i=1}^{k-1} \Delta p'(b_i) \right] U_x'(b_{n/2}) \iff U_x''(b_{n/2}) \leq U_x'(b_{n/2}) \\ &\iff b_{n/2} \text{ is moved closer to } x \end{aligned}$$

where the same statements hold when x is replaced by z . We can now moving the balls at the position of ball $b_{n/2}$ rightwards to x to create the final urns P' and Q' . When we do so, we are bringing all the balls closer to both x and z , which we just showed must decrease $\Delta U_x''$ and $\Delta U_z''$. Then $\Delta U_x' \leq \Delta U_x'' \leq U_x$ and $\Delta U_z' \leq \Delta U_z'' \leq U_z$, which concludes our proof for this case.

Case 3: $b_{n/2} > z$ for urn P The proof is symmetric to that of Case 2.

B Supporting Proofs

Theorem 8. [20] *Let a fixed size urn with R_0 red balls out of n total balls have an urn function $f(p)$ for which $\frac{f(p)}{f(1-p)}$ is monotonically decreasing and let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbb{E}[T] \leq \frac{1}{q_1} + \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \frac{q_k}{q_{k+1}(q_k - p_k)}$$

where $p_k = f\left(\frac{n-k}{n}\right)$ and $q_k = f\left(\frac{k}{n}\right)$.

Corollary 2. *Let a fixed size urn with R_0 red balls out of n total balls have an urn function $f(p) = 3p(1-p)^2$ and let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbb{E}[\tau] \leq n \ln n + O(n)$$

Proof. From Theorem 8, we know that

$$\mathbb{E}[T] \leq \frac{1}{q_1} + \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \frac{q_k}{q_{k+1}(q_k - p_k)} = \frac{n^3}{3(n-1)^2} + \frac{n^3}{3} \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \underbrace{\frac{n-k}{(k+1)(n-k-1)^2(n-2k)}}_{(*)}$$

where

$$\begin{aligned} (*) &= -\frac{1}{n(n-2)(n-k-1)^2} + \frac{4n}{(n+2)(n-2)^2(n-2k)} \\ &\quad + \frac{n+1}{n^2(n+2)(k+1)} - \frac{n^2+n-2}{n^2(n-2)^2(n-k-1)} \\ &\leq \frac{4n}{(n+2)(n-2)^2(n-2k)} + \frac{n+1}{n^2(n+2)(k+1)} \\ &= \left(\frac{4}{n^2(n-2k)} + \frac{1}{n^2(k+1)} \right) \left(1 + O\left(\frac{1}{n}\right) \right) \end{aligned}$$

Putting it together, we have

$$\mathbb{E}[T] \leq \frac{n}{3} + o(1) + \frac{n}{3} \left(1 + O\left(\frac{1}{n}\right) \right) \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \left[\frac{4}{n-2k} + \frac{1}{k+1} \right]$$

Noting that $\sum_{i=1}^k \frac{1}{k} = H_k = \ln n + O(1)$, where H_k is the k -th Harmonic number, we have

$$\begin{aligned} \mathbb{E}[T] &\leq O(n) + \frac{n}{3} \left(2 \ln \frac{n}{2} + \ln \frac{n}{2} \right) \\ &= n \ln n + O(n) \end{aligned}$$

Lemma 10. Let $f(x)$ be a concave monotonically decreasing function. Then for $s_j^1, s_j^2, t_i^1, t_i^2 \in \mathbb{R}$ satisfying $s_j^1 \leq t_i^1$, $s_j^2 \leq t_i^2$, $s_i^1 \leq s_j^2$, and $t_i^1 \leq t_i^2$ and $\text{sign}(c_i) = \text{sign}(d_j)$, we have

$$\frac{\sum_{i=1}^m c_i [f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^n d_j [f(s_j^2) - f(s_j^1)]} \geq \frac{\sum_{i=1}^m c_i [t_i^2 - t_i^1]}{\sum_{j=1}^n d_j [s_j^2 - s_j^1]} \quad \text{and} \quad \frac{\sum_{i=1}^m c_i [f(s_i^1) - f(s_i^2)]}{\sum_{j=1}^n d_j [f(t_j^1) - f(t_j^2)]} \leq \frac{\sum_{i=1}^m c_i [s_i^1 - s_i^2]}{\sum_{j=1}^n d_j [t_j^1 - t_j^2]}$$

Proof. Since f is concave, $s_j^1 \leq t_i^1$, and $s_j^2 \leq t_i^2$,

$$\frac{f(t_i^2) - f(t_i^1)}{t_i^2 - t_i^1} \leq \frac{f(t_i^2) - f(s_j^1)}{t_i^2 - s_j^1} \leq \frac{f(s_j^2) - f(s_j^1)}{s_j^2 - s_j^1}$$

Then by noting that f is monotonically decreasing, $s_j^1 \leq s_j^2$, $t_i^1 \leq t_i^2$, and $\text{sign}(c_i) = \text{sign}(d_j)$, we achieve the statement for a single term on the top and bottom

$$\frac{c_i [f(t_i^2) - f(t_i^1)]}{d_j [f(s_j^2) - f(s_j^1)]} \geq \frac{c_i [t_i^2 - t_i^1]}{d_j [s_j^2 - s_j^1]}$$

If we flip this inequality, we get

$$\frac{d_j[f(s_j^2) - f(s_j^1)]}{c_i[f(t_i^2) - f(t_i^1)]} \leq \frac{d_j[s_j^2 - s_j^1]}{c_i[t_i^2 - t_i^1]}$$

Using this, we can derive,

$$\frac{\sum_{j=1}^n d_j[f(s_j^2) - f(s_j^1)]}{c_i[f(t_i^2) - f(t_i^1)]} = \sum_{j=1}^n \frac{d_j[f(s_j^2) - f(s_j^1)]}{c_i[f(t_i^2) - f(t_i^1)]} \leq \sum_{j=1}^n \frac{d_j[s_j^2 - s_j^1]}{c_i[t_i^2 - t_i^1]} = \frac{\sum_{j=1}^n d_j[s_j^2 - s_j^1]}{c_i[t_i^2 - t_i^1]}$$

Then by flipping this inequality, we get

$$\frac{c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^n d_j[f(s_j^2) - f(s_j^1)]} \geq \frac{c_i[t_i^2 - t_i^1]}{\sum_{j=1}^n d_j[s_j^2 - s_j^1]}$$

Finally, we can use this to derive the first part of our final result

$$\frac{\sum_{i=1}^m c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^n d_j[f(s_j^2) - f(s_j^1)]} = \sum_{i=1}^m \frac{c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^n d_j[f(s_j^2) - f(s_j^1)]} \geq \sum_{i=1}^m \frac{c_i[t_i^2 - t_i^1]}{\sum_{j=1}^n d_j[s_j^2 - s_j^1]} = \frac{\sum_{i=1}^m c_i[t_i^2 - t_i^1]}{\sum_{j=1}^n d_j[s_j^2 - s_j^1]}$$

We can get the second part by simply inverting and multiplying the top and bottom of both sides by -1

$$\frac{\sum_{j=1}^n d_j[f(s_j^2) - f(s_j^1)]}{\sum_{i=1}^m c_i[f(t_i^2) - f(t_i^1)]} \leq \frac{\sum_{j=1}^n d_j[s_j^2 - s_j^1]}{\sum_{i=1}^m c_i[t_i^2 - t_i^1]}$$